


How To Recognize a Great Idea

To do ...

- Why and how to read



Outline

- Why reading
- Deciding what to read
- Suggestions on how to read

Need for efficient reading

- Why reading papers?
 - Review for a conference or a class
 - Literature survey of a new field
 - To keep up to date
- Too many papers!
 - SIGCOMM ~40, IMC ~42, ICDCS ~87, NSDI ~40, CoNEXT ~32, ...
 - Per year over 200
 - Hundreds of hours per year at it!

How and what to read

- A critical but rarely taught skill
- Not like reading fiction!
 - **Not** from start to end, but in passes
 - Not all at the same level, skip some/much
- A variant of the three-pass approach*
 - Get a bird's-eye view
 - Read as the audience
 - Read as an author

Deciding what to read

- *Goal:* To keep up to date
- *Approach:* Scan papers in latest conference

- *Goal:* To get up to speed on a subfield
- *Approach:* Transitive closure of related work of papers in top conference

- *Goal:* To increase your expertise
- *Approach:* Deep diving on key papers

Reading the news – Conf. proceedings

- Goal: To grasp main idea of a collection of papers; keep informed about problems and recent solutions
- Top-down method
 - Skim table of content: papers clustered into “sessions” which typically identify main areas
 - Consider authors and prioritize by area of interest and reputability of authors
 - As everything else, gets easier with time

How to read

- First pass – a few minutes
 - Title, abstract and introduction
 - Sections and subsection headings
 - Conclusions
 - Glance over references
- You should be able to answer
 - Category – measurement, analysis, ...
 - Context
 - Correctness – *Valid assumptions?*
 - Contributions
 - Clarity

How to read

- If you are interested in the paper but it's not your specialty or
- If you are exploring a new field ...
- Second pass ~1hr
 - Read with care, but gloss over details
 - Look at figures, diagrams
 - Mark relevant unread references for further reading
- At the end, you can summarize the main points, with supporting evidence to another person

Literature survey

- Goal: To get up to speed on a subfield
- Start by creating the seed
 - Recent papers from top conference, do a first pass
 - Survey paper if one exists
 - Seminal paper if different from above
- Else, perform transitive closure of cited work
 - Read related work sections of above papers
 - Focus on shared citations and repeated author names

Deep diving

- Goal: To understand some problem area in greater depth, on the path to be an expert
- Find the seminal paper in the field
- Read carefully, including evaluation
- Third pass ~1-5hr (depends on experience)
 - Try to virtually re-implement the paper
 - Identify and challenge every assumption
 - Think of how you would have presented the ideas
 - Comparing your virtual version with the actual paper provides a deeper insight

Questions to answer as you read

- Comprehension questions
 - You are reading to understand what the paper says
- Evaluation questions
 - But you have to be critical of the claims
- Synthesis
 - To generate research agendas

Invariant comprehension questions

- What is the problem?
 - Motivation? What is the niche?
- What are the claimed contributions?
 - What's new? Question, methodology, algorithm, research area ...

Invariant comprehension questions

- How do the authors substantiate the claims?
 - Methodology, argument, experimentation ... what makes it scientific instead of mere opinions?
- What are the conclusions?
 - What have you learned? Are the results generalizable? What are the open problems? ...
- *Where?* Abstract

Invariant evaluation questions

- Is the research problem significant?
- What is the “intellectual nugget”?
 - Every paper should have a key intellectual contribution
 - This key idea will also give your brain a way to “index” the paper

Invariant evaluation questions

- What are the main contributions?
 - New finding? Method? Perspective?
 - Are the contributions significant?
- Are the claims valid?
 - Have they cut corners? Error in proof? Problematic experimental setup? Are the claims modest enough?
 - Are the result likely to be affected by the method?

Keeping notes

- One sentence summaries are infinitely better than nothing at all
- Primitive approach – single file of notes
- Better – database with bibtex
 - There are some existing tools for bib management
 - Zotero, Papers, Mendeley, EndNote, ...
 - Will also help you more quickly construct related work sections for papers